Quasi-isometric embeddings of symmetric spaces

DAVID FISHER

KEVIN WHYTE

This paper opens the study of quasi-isometric embeddings of symmetric spaces. The main focus is on the case of equal and higher rank. In this context some expected rigidity survives, but some surprising examples also exist. In particular there exist quasi-isometric embeddings between spaces X and Y where there is no isometric embedding of X into Y. A key ingredient in our proofs of rigidity results is a direct generalization of the Mostow–Morse lemma in higher rank. Typically this lemma is replaced by the *quasiflat* theorem, which says that the maximal quasiflat is within bounded distance of a finite union of flats. We improve this by showing that the *quasiflat* is in fact flat off of a subset of codimension 2.

1 Introduction

The rigidity theorems of Mostow and Margulis are among the most celebrated results about the intersection of discrete groups and geometry. With the rise of Gromov's program [10] for the geometric study of discrete groups, coarse analogues of these results were among the most desired results. There are many possible translations of these theorems to a coarse setting, and several results and questions in this direction (see Farb [7] for a good survey). We first recall two basic definitions:

Definition 1.1 Let (X, d_X) and (Y, d_Y) be metric spaces. Given real numbers $K \geq 1$ and $C \geq 0$, a map $f \colon X \to Y$ is called a (K, C)–*quasi-isometry* if

(1) $\frac{1}{K} d_X(x_1, x_2) - C \leq d_Y(f(x_1), f(x_2)) \leq K d_X(x_1, x_2) + C$ for all x_1 and x_2 in X, and

(2) the C neighborhood of $f(X)$ is all of Y.

If f satisfies (1) but not (2), then f is called a (K, C)–*quasi-isometric embedding*.

Remark Throughout this paper, all symmetric spaces will have no compact or Euclidean factors.

To quickly summarize the current state of knowledge, good analogues of Mostow rigidity are now known in the coarse setting. In particular, from the work of many people, one can deduce:

Theorem 1.2 *Two symmetric spaces are quasi-isometric if and only if they are isometric (after a possible rescaling). Further, any finitely generated group quasi-isometric to a symmetric space is virtually a cocompact lattice in its isometry group.*

See Farb [7] for a detailed discussion and attribution. This theorem is not quite as strong as one would like to parallel Mostow rigidity — in particular, self-quasi-isometries of symmetric spaces or quasi-isometries between cocompact lattices in the same semisimple Lie group can be quite wild, while Mostow's result says group isomorphisms are induced by isometries. That is simply the truth for $\mathbb{H}^n$ and $\mathbb{C}\mathbb{H}^n$. For the other irreducible symmetric spaces, by results of Kleiner and Leeb [12] and Pansu [18] one has more:

Theorem 1.3 *Let X be an irreducible symmetric space of higher rank, or a quaternionic or Cayley hyperbolic space. Every quasi-isometry of X is at bounded distance from an isometry.*

This very satisfying analogue of Mostow's rigidity results can be used to give quick proofs of Mostow's results for cocompact lattices (coarse analogues are also known for lattices which are not cocompact, but those are generally more difficult; see Eskin [5]). Margulis's superrigidity results allow for distinct domain and range and homomorphisms rather than isomorphisms, and finding analogues of those results is an important problem in geometric group theory. The most obvious geometric question along these lines is to ask whether quasi-isometric embeddings of one symmetric space in another must be at bounded distance from the inclusion of a totally geodesic symmetric subspace, a question raised eg in [7].

In this paper we make the first significant progress on this question. Perhaps the most surprising piece of the puzzle is that exotic embeddings exist even assuming the domain and range have equal rank at least 2.

Theorem 1.4 *For any $r > 1$ there are quasi-isometric embeddings of the space* $\mathrm{SL}_{r+1}(\mathbb{R})/O(r+1)$ *into* $\mathrm{Sp}_{2r}(\mathbb{R})/U(2r)$.

The theorem gives quasi-isometric embeddings between spaces of rank r when there are no isometric embeddings in these cases, so the embeddings are definitely exotic. See Section 2 for a more detailed discussion of the behavior of these embeddings. Our

construction is more general, and produces quasi-isometric embeddings in most of the cases where our rigidity results fail. The heart of the construction is:

Theorem 1.5 *Let G_1 and G_2 be semisimple Lie groups of equal real rank with Iwasawa decompositions $G_i = K_i A_i N_i$. Every injective homomorphism of the solvable Lie group $A_1 N_1$ as a subgroup of $A_2 N_2$ is a quasi-isometric embedding.*

This is relevant to the above as every symmetric space $K \backslash G$ is isometric to the solvable Lie group AN coming from the Iwasawa decomposition $G = KAN$. We refer to quasi-isometries constructed using Iwasawa decompositions and homomorphisms as in Theorem 1.5 as AN–maps. Isometric embeddings, arising from homomorphisms from G_1 into G_2, are trivially also AN–maps. As another application of Theorem 1.5, we have:

Theorem 1.6 *There is a quasi-isometric embedding of $\mathbb{H}^2 \times \mathbb{H}^2$ into $\mathrm{SL}(3, \mathbb{R})/O(3)$. More generally, there is a quasi-isometric embedding of the product of n copies of $\mathbb{H}^2$ into $\mathrm{SL}(n + 1, \mathbb{R})/O(n + 1)$.*

It is not hard to adapt the proof of Theorem 1.6 to produce quasi-isometric embeddings of $\prod_{i=1}^{k} \mathrm{SL}(n_i + 1, \mathbb{R})$ into $\mathrm{SL}(n+1, \mathbb{R})$ whenever $\sum_{i=1}^{k} n_i = n$. We make no attempt to be complete in our discussion of quasi-isometric embeddings in the reducible case, focusing instead on a particular construction that shows a somewhat more dramatic failure of rigidity. By combining Theorems 1.6 and 1.4 with a simple *twist* described in Section 2 we prove the following theorem:

Theorem 1.7 *There exists a quasi-isometric embedding of the product of $\mathbb{H}^2 \times \mathbb{H}^2$ into $\mathrm{SP}(4, \mathbb{R})/U(2)$ that is not at bounded distance from a linear map along any flat and so is not at bounded distance from an AN–map.*

The behavior of these maps on flats is particularly surprising and we describe it in detail in Section 2. It is probably true that the same type of argument yields quasi-isometries that are not AN–maps from $(\mathbb{H}^2)^r$ to $\mathrm{Sp}(4r, \mathbb{R})/U(2r)$, but proving this by our method would involve some quite complicated combinatorics. While this paper was being revised, Nguyen [17] produced examples of quasi-isometric embeddings of $(\mathbb{H}^2)^n$ into $\mathrm{SL}(n + 1, \mathbb{C})$ where every flat maps over more than a single flat, and so the map is transparently not at bounded distance from an AN–map. We do not know whether all quasi-isometric embeddings between symmetric spaces of equal and higher rank are at bounded distance from compositions of AN–maps, although that seems unlikely.

In the other direction, despite the constructions in Theorems 1.4, 1.5, 1.6 and 1.7, there is a substantial amount of rigidity for quasi-isometric embeddings in equal and higher rank. We prove that quite often one has only isometric embeddings. The key issue seems to be linear embeddings of the patterns of hyperplanes in the restricted root system. We refer to this pattern of hyperplanes as the *Weyl pattern* of the symmetric space. The reason to believe that rigidity should occur in equal and higher rank is the *quasiflat theorem* of Eskin and Farb, and Kleiner and Leeb, which shows that the image of a flat in this setting is within a bounded distance of a finite union of flats. To prove our rigidity results, we require a substantial improvement to this statement, which we describe immediately after stating our main rigidity result.

Theorem 1.8 *Let X_1 and X_2 be irreducible symmetric spaces or Euclidean buildings, both of rank $r > 1$. Let C_1 and C_2 be their Weyl patterns on $\mathbb{R}^r$. Then:*

(1) *If there are no elements of $\mathrm{GL}_r(\mathbb{R})$ embedding C_1 into C_2, then there are no quasi-isometric embeddings of X_1 into X_2.*

(2) *Fix $K > 1$. If all embeddings of C_1 into C_2 which are K–quasiconformal are conformal, then all (K, C)–quasi-isometric embeddings of X_1 into X_2 are at bounded distance from a totally geodesic embedding. In particular, if all pattern embeddings are conformal then all quasi-isometric embeddings are at bounded distance from totally geodesic.*

Remark While part (1) of the conclusion of Theorem 1.8 follows easily for reducible symmetric spaces and buildings, the assumption of part (2) never holds in the reducible case, not even in the context of Corollary 1.11. Results in the reducible case have recently been obtained by Nguyen [17].

To explain the key technical mechanism in the proof, our version of the Mostow–Morse lemma, we pass to asymptotic cones and bilipschitz maps. We do so simply to have cleaner statements of the results. We now consider a building Y of rank k and in this context Kleiner and Leeb [12] showed that any bilipschitz map $f\colon \mathbb{R}^k \to Y$ has image contained in a finite union of flats. One can also deduce this by taking asymptotic cones of a theorem of Eskin and Farb [6]. Given such a map f, we call a point $x \in \mathbb{R}^k$ *flat* if there exists $r > 0$ such that $B(x, r) \subset \mathbb{R}^k$ maps to a single flat in Y. Given f, we will call the set of nonflat points for f the exceptional set for f.

Lemma 1.9 (higher-rank Mostow–Morse) *Let Y be an affine building of real rank k and $f\colon \mathbb{R}^k \to Y$ a bilipschitz map. Then the exceptional set Z for f has codimension 2.*

In fact, Z is contained in a finite collection of bilipschitz images of linear spaces of dimension $k-2$.

Since codimension-two sets in one-dimensional flats are empty, this statement is a direct generalization of the standard Mostow–Morse lemma in hyperbolic spaces. This lemma is proven in Lemma 3.1 and is a necessary component of our proofs. While we do not investigate this here, it seems that Lemma 1.9 should be true in a much broader context, namely any setting where there is (a) a quasiflat theorem, and (b) some combinatorial control on intersections of flats, including perhaps settings like those of Behrstock, Hagen and Sisto [1] and Huang [11]. We prove the lemma from the quasiflats theorem; it seems quite difficult to prove directly without first proving that theorem. There is a coarse analogue of Lemma 1.9 for quasiflats in symmetric spaces and buildings, but the coarse version of the statement becomes quite cumbersome.

We use Lemma 1.9 in proving both parts of Theorem 1.8, though a proof of the first part is possible, though much more cumbersome, without it. To prove both parts of the theorem, we pass to asymptotic cones and study the map restricted to flats. To prove the first part of the theorem, we look at a flat point x of a flat F where we have a bilipschitz map of domains in $\mathbb{R}^r$. By Rademacher and Stepanov this map is differentiable almost everywhere and the goal is to see that the derivative is a linear embedding of patterns. To check this, one uses that the pattern arises as the intersection of F with transverse flats to control the derivative. To prove this one uses that these intersections are of codimension one in both flats and so are not contained in the exceptional sets. The proof of the second conclusion is more difficult. The key step is to show that a flat maps to a single flat. The hypothesis allows us to see that the derivative along a flat is everywhere conformal and so the map on the flat is smooth off the exceptional set by a result of Gehring. The hypotheses imply that the isometric part of the derivative is an element of the Weyl group at every point and the fact that the exceptional set is of codimension two allows us to conclude this Weyl group element is constant and then check that the rescaling must be the same at every point. This shows the map on the flat is globally isometric.

The use of Lemma 1.9 leads to a short elegant proof in our context and also allows for significant simplification of the proofs of Druţu [4] and Eskin [5]. To conclude the proof of Theorem 1.8 in the setting of symmetric spaces requires an additional input, namely the main result of Kleiner and Leeb [13], which we use to conclude that the image of the quasi-isometric embedding, which is a union of isometrically embedded flats, is actually a subsymmetric space.

We now state some corollaries of Theorem 1.8. In Section 4 we classify all pattern embeddings (note that B_n, C_n and BC_n all have the same pattern, which we will call BC_n). The conclusion is that these are all conformal except for some exceptional embeddings of A_n into BC_n or F_4 and some exceptional embeddings of A_2 and BC_2 into G_2.

Corollary 1.10 *Among equal-rank, irreducible buildings or symmetric spaces, all quasi-isometric embeddings are at bounded distance from totally geodesic unless the domain is of type A_n and the range of type B_n, C_n, BC_n or F_4 or unless the source is of type D_4 and the range is of type A_4 or unless the range is of type G_2.*

A simple example of this is that all quasi-isometric embeddings of $\mathrm{SL}_n(\mathbb{R})$ into $\mathrm{SL}_n(\mathbb{C})$ (for $n > 2$) are near-isometric embeddings. As discussed in Sections 2 and 4, our rigidity results and construction are close to complementary. See Section 5 for a discussion of the few cases where neither applies. We also have an additional corollary concerning quasi-isometric embeddings with "small constants".

Corollary 1.11 *Let X and Y be irreducible buildings or symmetric spaces of equal rank $r > 1$. Then there exists a constant K, depending only on the Weyl patterns of X and Y, such that for any $K' < K$ and any C, any (K, C)–quasi-isometric embedding of X into Y is at bounded distance from an isometry.*

Remark One might expect the gap in the corollary to be a result of a limiting argument, taking (K_i, C)–quasi-isometric embeddings for some sequence $K_i \to 1$ and then analyzing the case of $(1, C)$–embeddings. That $(1, C)$–embeddings are necessarily isometric can be proven using a somewhat simplified version of the proof of Theorem 1.8. This limiting argument would yield only some nonexplicit K slightly larger than 1. If one proves the corollary directly from Theorem 1.8 instead, one obtains an explicitly computable K that depends only on the Weyl patterns of X and Y.

The ideas used in this paper find further application in work of the first author and Nguyen [8] on quasi-isometric embeddings of nonuniform higher-rank lattices.

We end the introduction with some remarks on the case of embeddings between rank-one spaces and embeddings where the rank is allowed to increase. Without the assumptions of equal rank and higher rank, the picture becomes quite opaque and it is not clear that any rigidity remains. In rank one, a theorem of Bonk and Schramm [2] shows, in particular, that every rank-one symmetric space quasi-isometrically embeds in $\mathbb{H}^n$ for n

sufficiently large (and hence into $\mathbb{CH}^n$ and $\mathbb{HH}^n$). Thus, quasi-isometric embeddings exist in many settings where there are no isometric embeddings. One can iterate these constructions and combine them with isometric embeddings to build various exotic quasi-isometric embeddings between, say, $\mathbb{HH}^n$ and $\mathbb{HH}^m$ for some $m \gg n$. Further, even when isometric embeddings do exist, quasi-isometric embeddings can be quite far from isometric, even equivariantly; for example, quasifuchsian groups in $\mathbb{H}^3$ and more general bendings of cocompact lattices in $\mathbb{H}^n$ inside of $\text{Isom}(\mathbb{H}^{n+1})$ are quasi-isometric as long as they remain convex cocompact.

When the rank of the target is higher than that of the domain things are also not well controlled, even if both are of higher rank (since rank increases under quasi-isometric embeddings, there are none where the rank of the domain exceeds that of the target). For example, any X can be quasi-isometrically embedded in $X \times \mathbb{R}$ as the graph of any Lipschitz function $X \to \mathbb{R}$. In particular, there are many strange embeddings of $\text{SL}(n, \mathbb{R})$ into $\text{SL}(n + 1, \mathbb{R})$ of this type as $\text{SL}(n, \mathbb{R}) \times \mathbb{R}$ is a totally geodesic (and full-rank) subspace of $\text{SL}(n + 1, \mathbb{R})$.

In both of these cases it still seems to be a difficult question to determine precisely when quasi-isometric embeddings exist; see Section 5 for some open problems along these lines.

2 AN–maps

In this section we describe a construction of some "exotic" quasi-isometric embeddings. These examples are built from injective homomorphisms of solvable Lie groups. Let G be a semisimple Lie group, with Iwasawa decomposition $G = KAN$. The solvable group AN acts simply transitively on $X = K \backslash G$ by isometries, and so is quasi-isometric to X. The key observation in our construction is:

Proposition 2.1 Let $A = \mathbb{R}^k$ for some $k \geq 1$, and let N_1 and N_2 be two nilpotent Lie groups with A–actions such that N_i is uniformly exponentially distorted in AN_i. If $f \colon N_1 \to N_2$ is an A–equivariant embedding of Lie groups then the induced embedding $AN_1 \to AN_2$ is a quasi-isometric embedding.

Remark We call maps satisfying the hypotheses of the proposition AN–maps.

Proof We use the notation $\|g\|_G$ throughout for the distance in G from g to the identity. Here uniformly exponentially distorted means that there is a $C > 0$ such that for all $n \in N$ we have $\|n\|_{AN} \leq C \log(1 + \|n\|_N)$. This holds for the solvable groups AN coming from semisimple Lie groups as N is spanned by nontrivial root spaces. It holds in many other, but not all, solvable Lie groups.

Let λ be such that f is λ–Lipschitz (such a constant exists as the map $AN_1 \to AN_2$ is a homomorphism). The content of the proposition is that there is a linear lower bound on distances in AN_2 in terms of the distances in AN_1. Since f is a homomorphism, it suffices to find such a bound for the distance from the identity.

The following holds because the exponential map for nilpotent groups is polynomial.

Claim 2.2 Let N and N' be simply connected nilpotent Lie groups. For any embedding N as a closed subgroup of N', there is a polynomial P such that for all $n \in N$ we have $\|n\|_N \leq P(\|n\|_{N'})$.

We now return to the proof of the proposition. Let $g = an$ be an element of AN as above. The triangle inequality gives $\|g\| \leq \|n\| + \|a\|$. Note that all norms appearing in this argument are norms in the group AN, not in any subgroup. Since the projection to A is distance-decreasing, $\|g\| \geq \|a\|$. We also have

$$\|g\| = d(an, \mathrm{Id}) \geq d(an, a) - d(a, \mathrm{Id}) = \|n\| - \|a\|.$$

Thus, if $\|a\| \leq \frac{1}{2}\|n\|$, we have $\|g\| \geq \frac{1}{2}\|n\|$. Otherwise, $\|g\| \geq \|a\|$. So we have

$$\|a\| + \|n\| \geq \|g\| \geq \tfrac{1}{2}\max(\|a\|, \|n\|) \geq \tfrac{1}{4}(\|a\| + \|n\|).$$

Thus, the distance to the identity is bilipschitz-equivalent to the distance in the product $A \times N$, where A and N are given the norms induced by their inclusions in AN. Since f is a linear isomorphism along A, it is bilipschitz there. For f embedding N_1 into N_2 we have that the norm in N_2 is bounded by a polynomial in the norm in N_1 by the claim, and so, by uniform exponential distortion, the norm of an n in AN_2 is bounded by the logarithm of a polynomial of an exponential of the norm in AN_1. This yields the desired linear lower bound. $\qquad\square$

For $G = KAN$ we have that N splits as a sum of positive root spaces $N = \bigoplus_\lambda E_\lambda$, where λ is a linear function on A describing the action of A on the part of the Lie algebra of N tangent to E_λ. An A–equivariant embedding of N_1 into N_2 must send these root spaces into root spaces.

If λ_1 and λ_2 are positive roots in G_1, then either $\lambda_1 + \lambda_2 = \lambda$ for a root λ, in which case $[E_{\lambda_1}, E_{\lambda_2}] = E_\lambda$, or $\lambda_1 + \lambda_2$ is not a root, in which case E_{λ_1} and E_{λ_2} commute. Since a homomorphism respects brackets, it follows that if E_{λ_i} maps into the E_{η_i} and $\lambda_1 + \lambda_2 = \lambda$ then $\eta_1 + \eta_2 = \eta$ for a root η and E_λ maps into the η–root space of N_2.

When $\lambda_1 + \lambda_2$ is not a root in G_1, we must have that the images in the η_1 and η_2 root spaces in G_2 commute. For simplicity, assume that G_2 is $\mathbb{R}$–split or complex. This then implies that $\eta_1 + \eta_2$ is not a root in G_2. Under this assumption, we have an additive map $\lambda \mapsto \eta$ sending positive roots for G_1 to positive roots for G_2 such that two roots in the domain sum to a root if and only if they do in the range. This linear map is nothing but the map induced on A. Conversely, given such a map, one clearly has an embedding of solvable groups. In summary:

Proposition 2.3 *Let G_1 and G_2 be semisimple real Lie groups with restricted root systems R_1 and R_2 of equal rank. There is an embedding of AN_1 into AN_2 if there is a linear automorphism T of A carrying positive roots into positive roots and such that λ_1 and λ_2 sum to a root if and only if $T(\lambda_1)$ and $T(\lambda_2)$ do. If the G_i are both $\mathbb{R}$–split, both complex, or G_1 is split and G_2 is complex, then the converse holds.*

Corollary 2.4 *There is a quasi-isometric embedding of $\mathrm{Sl}_{n+1}(\mathbb{R})$ into $\mathrm{Sp}_{2n}(\mathbb{R})$.*

Proof By the previous proposition, we must find a linear map $\mathbb{R}^n \to \mathbb{R}^n$ which sends the positive roots of A_n to positive roots of C_n respecting when roots sum to roots. Represent the root systems as

- A_n with positive roots $x_i - x_j$ for $i > j$ on the space,
- $\mathbb{R}^n = \{(x_0, \ldots, x_n) : x_0 + \cdots x_n = 0\}$, and
- C_n with positive roots $2y_i$, $y_i + y_j$ and $y_i - y_j$ for $i > j$ on $\mathbb{R}^n = \{(y_1, \ldots, y_n)\}$;

then the map

$$T(x_0, \ldots x_n) = \tfrac{1}{2}(2x_1 - (x_0 + x_n), 2x_2 - (x_0 + x_n), \ldots, 2x_{n-1} - (x_0 + x_n), x_n - x_0)$$

sends positive roots to positive roots; the root for $x_i - x_j$ for $i > j > 0$ corresponds to the root $y_i - y_j$, the root $x_i - x_0$ for $0 < i < n$ to $y_i + y_n$, and $x_n - x_0$ to $2y_n$.

All that remains is to verify that this correspondence respects which pairs of positive roots sum to a root. This is straightforward for this particular map but can be subtle in greater generality; see Section 4.2 and particularly Lemma 4.13. □

Remark This same map gives a quasi-isometric embedding $\mathrm{Sl}_{n+1}(\mathbb{C})$ into $\mathrm{Sp}_{2n}(\mathbb{C})$, or any other A_n to C_n situation where the domain is $\mathbb{R}$–split (or complex if the target is as well). The exact same construction works over p–adic fields. It is interesting to compare this to the linear pattern embeddings classified in Section 4.

Remark There is some earlier work using AN–maps to build quasi-isometric embeddings, but only in cases where A is one-dimensional. In particular, Brady and Farb [3] use a similar construction to show that $\mathbb{H}^3$ quasi-isometrically embeds in $\mathbb{H}^2 \times \mathbb{H}^2$. In rank one, this construction is quite flexible and has other applications [3; 9; 15], but it remains quite surprising that such a construction is possible in the more combinatorially complex higher-rank case.

It is interesting to look at the geometry of this map in more detail. Restricted to a flat in Sl_{n+1} corresponding to a coset of A in AN, the map f is precisely the linear map T above. The hyperplanes in Sp_{2n} that do not come from those in Sl_{n+1} come from the roots $2y_i$ for $i < n$ and $y_i + y_j$ for $i < j < n$. These correspond to the functions $2x_i = x_0 + x_n$ (for $0 < i < n$) and $x_i + x_j = x_0 + x_n$ (for $0 < i < j < n$). The chambers of the flat in the domain are given by the ordering of the coordinates, and how these missing hyperplanes subdivide such a chamber depends on where x_0 and x_n sit in this ordering.

A flat in A_n has $(n + 1)!$ chambers and one in C_n has $2^n n!$. Thus, the "average" chamber must map across $2^n / (n + 1)$ chambers. The minimally subdivided chambers are those where x_0 and x_n are either the two smallest or two largest in the ordering (these come in four families from the ordering between x_0 and x_n and whether they are minimal or maximal) and map to single chambers in C_n. Those maximally subdivided are those with x_0 and x_n the minimum and maximum (in either order).

For concreteness we look at $n = 2$. There are six chambers in our flat, corresponding to the permutations of $\{0, 1, 2\}$. There are two chambers in which x_0 and x_2 are the minimum and maximum. The ordering $x_0 < x_1 < x_2$ is the unique chamber fixed by N (call this chamber C_0) and the ordering $x_2 < x_1 < x_0$ is the opposite chamber. Both of these are sent to two chambers in $\mathrm{Sp}_4 / U(2)$. The four orderings in which x_1 is either minimal or maximal are all sent to single chambers.

This gives a picture of what happens for an arbitrary flat. Divide the flat into its six chambers, and label each according to whether it is the chamber fixed by N (there is just one such), adjacent to this chamber, opposite to this chamber (the generic case), or adjacent to an opposite chamber. What happens to every chamber is then determined as above. It is not hard to see that the flats in Sl_3 are of the following types:

- A flat passing through C_0. These are described above, and have two chambers (opposite to each other in the flat) that map to two chambers, and the four others, which map to single chambers. Thus, the flat maps to a total of eight chambers which is necessarily a single flat. The map on flats is given by the linear map T above.

- There are two chambers of F (next to each other in F) which are opposite to C_0, their neighbors in F, which are adjacent to chambers opposite to C_0, and two chambers adjacent to C_0. This maps across eight total chambers, and hence has image at bounded distance from a single flat. Unlike the previous case, however, the map is not globally linear as a map $\mathbb{R}^2 \to \mathbb{R}^2$.

- There are two chambers of F (next to each other in F) which are both adjacent to C_0 or both adjacent to chambers opposite to C_0. The remaining four chambers are all opposite to C_0. Such a flat maps to an image quasiflat with ten chambers, and hence is not near a single flat.

- All chambers in F are opposite to C_0 (the generic case). In this case all six chambers map across two chambers in the range. The image quasiflat has twelve chambers and so is not near a single flat.

We now prove Theorem 1.6, which we can do by providing an explicit map on matrices. The Iwasawa decomposition of $(\mathbb{H}^2)^n$ can be realized as $AN = \prod A_i N_i$, where each A_i is a two-by-two diagonal matrix

$$\begin{bmatrix} a_i & 0 \\ 0 & a_i^{-1} \end{bmatrix}$$

and each N_i is a two-by-two diagonal matrix

$$\begin{bmatrix} 1 & u_i \\ 0 & 1 \end{bmatrix},$$

where each a_i is a nonzero real number and each u_i is a real number.

The desired AN–embedding is then explicitly

$$\begin{bmatrix} \alpha_1 & u_1 & u_2 & \cdots & & u_n \\ 0 & \alpha_2 & 0 & \cdots & & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 & \alpha_1{}^{-1}\cdots\alpha_n{}^{-1} \end{bmatrix},$$

where the α_i are easily computed explicitly. For instance, in the case $n = 2$, we have $\alpha_1 = a_1^2 a_2^2$ and $\alpha_2 = a_2^2$.

We examine the behavior of these maps, particularly how they behave along flats, in more detail. This leads to a proof of Theorem 1.7. First note that in an AN–map, every flat which contains the chamber at infinity fixed by AN maps linearly to a flat in the range. In particular, we will show the map we construct is not an AN–map by showing this does not occur (and, in fact, that no flat maps linearly from domain to range). We begin by describing the image of flats in $\mathbb{H}^2 \times \mathbb{H}^2$ into $SL(3, \mathbb{R})/O(3)$ and then combining this with the description given above for flats $SL(3, \mathbb{R})/O(3)$ mapping into $SP(4, \mathbb{R})/U(2)$.

For this description, we will think of $\mathbb{H}^2$ as being metrics on $\mathbb{R}^2$ up to scalars, and $SL_3(\mathbb{R})/SO_3(\mathbb{R})$ as the same on $\mathbb{R}^3$. The embedding is then: Choose two planes P_1 and P_2 in $\mathbb{R}^3$, and identify each with $\mathbb{R}^2$. A point of $\mathbb{H}^2 \times \mathbb{H}^2$ can then be thought of as a metric (up to scalar) on each of P_1 and P_2. Rescale so they agree on their line of intersection (let's call that l), and extend that to an inner product on $\mathbb{R}^3$ by having P_1 and P_2 meet orthogonally.

The geodesics in $\mathbb{H}^2$ correspond to the set of metrics on $\mathbb{R}^2$ for which some fixed pair of lines are orthogonal, and a maximal flat in $SL_3(\mathbb{R})/SO_3(\mathbb{R})$ corresponds to the set of metrics in which some fixed triple of lines are orthogonal. Choose lines $l_1 \subset P_1$ and $l_2 \subset P_2$, neither equal to l, and consider the flat in $\mathbb{H}^2 \times \mathbb{H}^2$ in which l, l_1 and l_2 are orthogonal. These are precisely the flats which pass through the chamber fixed by AN, and each maps linearly to the flat in $SL_3(\mathbb{R})/SO_3(\mathbb{R})$ in which l, l_1 and l_2 are orthogonal.

This flat in $\mathbb{H}^2 \times \mathbb{H}^2$ is subdivided into four quadrants — if we choose vectors $v \in l$, $v_1 \in l_1$ and $v_2 \in l_2$ then we can think of the quadrants as determined by which of v and v_i is longer in the set of metrics on P_i. The flat in $SL_3(\mathbb{R})/SO_3(\mathbb{R})$ is similarly divided into six chambers according to the ordering of sizes of v, v_1 and v_2. Thus, some chambers in the domain map to two chambers in the range and others to only one. For example, if we know that v is shorter than v_1 in P_1 and v is longer than v_2

in P_2, then, once we rescale to get the metrics to agree on v, we know that the order is $\|v_2\| < \|v\| < \|v_1\|$ and we have a quadrant that maps to a single chamber. On the other hand, if v is shorter than v_1 in P_1 and shorter than v_2 in P_2, after rescaling we might have $\|v\| < \|v_1\| < \|v_2\|$ or $\|v\| < \|v_2\| < \|v_1\|$, so this quadrant gets mapped across two chambers.

Viewing the chambers in $\mathbb{H}^2 \times \mathbb{H}^2$ as points in $\partial\mathbb{H}^2 \times \partial\mathbb{H}^2$, what we see is that a chamber in which neither point is l is mapped to two chambers, as is the one chamber with both points l, while those in which one is l and one is not are mapped to a single chamber. Thinking of chambers in $\mathrm{SL}_3(\mathbb{R})/\mathrm{SO}_3(\mathbb{R})$ as flags $V_1 \subset V_2 \subset \mathbb{R}^3$ and the points at infinity in $\mathbb{H}^2$ as lines in $\mathbb{R}^2$, what we have is this:

- A chamber in $\mathbb{H}^2 \times \mathbb{H}^2$ with points at infinity l_1 and l_2, neither equal to l, maps to the two chambers $l_1 \subset V$ and $l_2 \subset V$, where V is the span of l_1 and l_2.

- A chamber in $\mathbb{H}^2 \times \mathbb{H}^2$ with points at infinity l_1 and l, with $l_1 \neq l$, maps to the chamber $l_1 \subset P_1$. Likewise, a chamber with endpoints l and l_2 maps to the one corresponding to the flag $l_2 \subset P_2$.

- The chamber with both points at infinity equal to l maps to the two chambers $l \subset P_1$ and $l \subset P_2$.

From this it is easy to determine what happens to an arbitrary flat $F = \gamma_1 \times \gamma_2$ in $\mathbb{H}^2 \times \mathbb{H}^2$. If none of the endpoints of γ_i is l then the image is a quasiflat made of eight chambers; if both geodesics have l as an endpoint, the image is a single flat (and the map is linear here). Finally, if one of the geodesics has l as an endpoint and the other does not, the image is again a single flat (the metrics for which the three endpoints that are not l are orthogonal) but the map is not linear — one half of the flat maps to four of the chambers and the other half to two.

The embedding of $\mathrm{SL}_3(\mathbb{R})/\mathrm{SO}_3(\mathbb{R})$ into $\mathrm{Sp}_4(\mathbb{R})/U(2)$ can be described similarly (by extending a metric on $\mathbb{R}^3$ to one on $\mathbb{R}^4$ compatible with the standard symplectic form) but is more complicated. The summary for chambers follows as described above. In this terminology, there is a distinguished flag $V_1 \subset V_2$ in $\mathbb{R}^3$, and this flag as well as those opposite to it (meaning those $U_1 \subset U_2$ for which $U_1 \not\subset V_2$ and $V_1 \not\subset U_2$) map to two chambers. Flags neither equal to nor opposite to $V_1 \subset V_2$ map to a single chamber.

We can compose AN–maps to get an embedding of $\mathbb{H}^2 \times \mathbb{H}^2$ into $\mathrm{SP}_4(\mathbb{R})/U(2)$. The maps $\iota_1 \colon \mathbb{H}^2 \times \mathbb{H}^2 \to \mathrm{SL}(3, \mathbb{R})/O(3)$ and $\iota_2 \colon \mathrm{SL}(3, \mathbb{R})/O(3) \to \mathrm{SP}(4, \mathbb{R})/U(2)$ are each constructed using a choice of Iwasawa decomposition for $\mathrm{SL}(3, \mathbb{R})$. The

composition is only obviously an AN–map if we choose the same decomposition for both maps. We want to consider the other extreme, where the flags defining the two Iwasawa decompositions are opposite to one another. In the notation above, this means V_1 does not lie in either P_i and l does not lie in V_2. There are then two special lines, namely $l_1 = P_1 \cap V_2$ and $l_2 = P_2 \cap V_2$. Combining our calculations of the behavior of flats under AN–maps gives the following for the composition:

- The chamber in $\mathbb{H}^2 \times \mathbb{H}^2$ with endpoints (l_1, l_2) maps to the two chambers corresponding to the flags $l_1 \subset V_2$ and $l_2 \subset V_2$. These two are both adjacent to $V_1 \subset V_2$ and so each maps to a single chamber. The image under the composition thus is two chambers.

- The chamber (l_1, l) maps to the single flag $l_1 \subset P_1$. This is an adjacent chamber and so the image under the composition is a single chamber. The same holds for the symmetric case (l, l_2).

- The chamber (l, l) maps to the two flags $l \subset P_1$ and $l \subset P_2$. Both of these are opposite, and so the image under the composition is four chambers.

- A chamber (l_1, l_2') where $l_2' \notin \{l, l_2\}$ maps to the two flags $l_1 \subset U$ and $l_2' \subset U$, for U the span of l_1 and l_2'. The first is adjacent and the second opposite, so the image has three chambers. The same holds for the symmetric cases (l_1', l_2).

- A chamber (l, l_2') (with l_2' as above) maps to the single flag $l_2' \subset P_2$. This is opposite and so maps to two chambers. The same holds for the symmetric cases (l_1', l).

- There is a bijection between the lines in $P_1 \setminus \{l_1, l\}$ and lines in $P_2 \setminus \{l_2, l\}$ defined by saying two lines correspond if and only if their span contains V_1. A chamber with endpoints (l_1', l_2') maps to the two flags $l_1' \subset U$ and $l_2' \subset U$, where U is the span of l_1' and l_2'. These are both adjacent if l_1' and l_2' correspond under the above bijection, and otherwise are both opposite. In the former case the image under the composition is two chambers, and in the latter is four.

We can deduce that this map is not an AN–map from this: Fixing any chamber at infinity in $\mathbb{H}^2 \times \mathbb{H}^2$ we can find a flat passing through it which has the opposite chamber of the generic type (the second part of the final case above), and so has a quadrant which maps to a half flat. This shows the map along this map cannot be a linear map to any flat. This proves Theorem 1.7. A more careful examination of cases shows that, in fact, no flat maps linearly to a flat. If F were such a flat, we would have that every half-flat of F would map to a half-flat. This can only happen in two ways: if every

quadrant maps to two chambers, or if two opposite quadrants map to a single chamber and the other two to three chambers each. The second possibility is easily ruled out as the only two quadrants which map to single chambers are (l_1, l) and (l, l_2), so the only flat containing both is $(l_1, l) \times (l, l_2)$. This flat maps to a single flat, but not linearly as it sends the quadrant (l, l) to a half-flat. Similarly, we have that the quadrants mapping to two chambers are (l_1, l_2), those of the form (l, l_2') or (l_1', l), or of the form (l_1', l_2'), where l_1' and l_2' correspond under the bijection described above, and no flat in $\mathbb{H}^2 \times \mathbb{H}^2$ has all four quadrants of these forms, so no flat has every quadrant mapping to two chambers.

Remark In these examples there are some flats which map to flats, just not linearly. Nguyen [17] gives some examples quite similar to these of embeddings of $\mathbb{H}^2 \times \mathbb{H}^2$ into the symmetric space X for $\mathrm{SL}(3, \mathbb{C})$, where it is clear that every chamber in $\mathbb{H}^2 \times \mathbb{H}^2$ maps to two chambers in X and so every flat maps to eight chambers in X and so the map is not AN since no flat maps to a single flat. It would be interesting to be able to show maps are not AN without needing to analyze behavior along every flat.

3 Rigidity

If X_1 quasi-isometrically embeds in X_2 of equal rank, then every maximal flat in X_1 gives a quasiflat in X_2. By results of [12] and [6], such a quasiflat is at bounded distance from a finite union of flats. As we want to control the intersection of two quasiflats, we need more information on precisely what a quasiflat can look like. For simplicity we describe the results for bilipschitz embeddings of $\mathbb{R}^n$ into Euclidean buildings, since the results for quasiflats in Euclidean buildings or symmetric spaces follow formally by passage to asymptotic cones. For background on buildings, symmetric spaces and asymptotic cones and their relations to one another, we refer the reader to [12].

3.1 Structure of quasiflats in buildings

In this subsection, X will also be a Euclidean building of rank d. By a *flat* we mean a top-dimensional, isometrically embedded $\mathbb{R}^d$ in X. Such a flat comes with a distinguished family of codimension-one affine subspaces, which we call *Weyl hyperplanes*. An affine subspace of a flat which is the intersection of the hyperplanes containing it is called a *subflat*. Each subflat comes with an induced pattern of hyperplanes by intersection, and we continue to call these Weyl hyperplanes (of the subflat). By a *Weyl*

box in a (sub)flat we mean a convex set which is a finite intersection of closed Weyl halfspaces. We now state a more precise form of Lemma 1.9 from the introduction. The description of the exceptional set here takes place in the range rather than the domain, but the two are, of course, related by a bilipschitz map.

Lemma 3.1 *For any K, there is an n such that for any K–bilipschitz embedding f of $\mathbb{R}^d$ into X there is a finite set $\{S_1, S_2, \ldots, S_n\}$ of codimension-two subflats of X such that every point x in the image of f not contained in any of the S_i has a neighborhood contained in a flat.*

Proof   We begin by discussing the local structure of a union of two flats. For this we need some vocabulary. Given a finite collection of disjoint rays, we form a *star* by identifying their initial points. A *fan* is a product $F = S \times \mathbb{R}^k$ where S is a star. If S consists of exactly four rays, then F is a union of two flats intersecting transversely.

Lemma 3.2   *There is a k (depending only on X) such that for any flats F and F' of X, off of at most k codimension-two subflats, each point of $F \cup F'$ has a neighborhood which is either contained entirely in a flat or which is contained in a fan.*

Proof   The intersection $F \cap F'$ is a Weyl box in a (sub)flat of each. If this subflat is of codimension two (or more) then we are done (with $k = 1$). If it is of codimension one then the intersection is a box in a hyperplane H. The interior points of the box have neighborhoods (in $F \cup F'$) which are the contained in the union of two flats which forms a fan. Thus, the points at which the conclusion of lemma fails are the boundary points of the intersection. The number of faces of such a box is bounded by the combinatorics of the Coxeter system of X, and each face is contained in a codimension-two subflat of F and F'.

Finally, if the intersection is of codimension zero, then the conclusion of the lemma holds in the interior and exterior of the intersection (these points are locally contained in a single flat) and at points in the interior of faces of the box of intersection (these are locally in a pair of transverse flats forming a fan). Thus, the points at which it fails are contained in the codimension-two and higher facets of the boundary of the box. As before, the number of these facets is bounded by the Coxeter system, and each facet is contained in a codimension-two subflat.   □

Corollary 3.3   *For any j there is an m such that for any collection $F_1, F_2, \ldots, F_j$ of flats, off of at most m codimension-two subflats (called **exceptional**), each point*

of $\bigcup F_i$ has a neighborhood which is either contained entirely in a flat or which is contained in the union of finite flats meeting pairwise transversely along a common hyperplane H.

Proof Let x be a point of the union $\bigcup F_i$ off of the exceptional sets produced by the previous lemma for the pairwise intersections. If the various hyperplanes through x along which the flats branch are the same in a neighborhood of x then we are done. If two of the flats intersect in something of codimension two or more then we add that subflat to the collection of exceptional subflats $\big($the number of such is bounded by $\binom{j}{2}\big)$. Thus, we need only address points at which three or more flats intersect.

If three flats intersect pairwise in a subflat of codimension two more then we add this intersection to the set of exceptional subflats.

Finally, if a three flats intersect near x in a box in a hyperplane H then the lemma holds near x unless x is on the boundary of the box, and is thus in one of the codimension-two subspaces containing the faces of the box. The number of these is similarly bounded by j and X. $\qquad\square$

We can now finish the proof of Lemma 3.1. By [12, Corollary 7.2.4], there is an s depending only on K such that the image of f is a subset of the union of at most s flats. Let Y be the union of these flats. Applying the corollary gives a family of exceptional subflats off of which every point of Y has a neighborhood which is either flat (in which case there is nothing to prove) or a finite collection of Weyl halfspaces meeting along a common hyperplane H. In the latter case, we locally have an embedding of $\mathbb{R}^d$ into $\mathbb{R}^{d-1} \times S$, where S is a finite union of rays meeting at a single point. As such an embedding must visit only two of the rays in S, we have the image contained in a union of two half spaces meeting along a hyperplane. Such a union is itself a flat in X. $\qquad\square$

3.2 Flats go to flats

In this subsection, we prove that under hypotheses analogous to those of Theorem 1.8, bilipschitz embeddings of Euclidean buildings take flats to flats.

Proposition 3.4 *Let X and Y be Euclidean buildings of equal rank such that every K–quasiconformal linear embedding of patterns is conformal. Then every K–bilipschitz embedding of X into Y sends flats to flats and is a similarity along every flat.*

Lemma 3.5 *Let ϕ be an embedding as above and F a flat in X. For any hyperplane H of F, almost every point $x \in H$ has a neighborhood mapped to a hyperplane of Y.*

Proof By Lemma 3.1, we know $\phi(F)$ is locally contained in a flat off of a finite set of codimension-two subflats. Let $\Sigma \subset F$ be the preimage of these exceptional subflats. For any $x \in F \setminus \Sigma$, the embedding ϕ is locally a map into a single flat. By Rademacher's theorem, ϕ is differentiable at almost every point in such a neighborhood, with derivative a K–quasiconformal linear map. Since these neighborhoods cover $F \setminus \Sigma$, this holds at almost every point of $F \setminus \Sigma$.

Let x be such a point of differentiability and H a hyperplane of F through x, let $y = \phi(x)$ and G be a flat in Y which locally contains $\phi(F)$ near y. Choose a flat F' in X which intersects F in H. The image of F' is also a bilipschitz image of a flat, and so has the structure given by Lemma 3.1. We have $\phi(H) = \phi(F) \cap \phi(F')$. Off of the exceptional points of $\phi(F')$ this is locally the intersection of two flats which are transverse as F and F' are. Thus, $\phi(H)$ is (locally) contained in the closure of a finite union of hyperplanes in G. By the differentiability of ϕ at x, $\phi(H)$ must locally be equal to a single such hyperplane. $\qquad\square$

Proof of Proposition 3.4 We now finish the proof of Proposition 3.4. As in the proof of the lemma, let x be a point of differentiability of ϕ in $F \setminus \Sigma$. Lemma 3.5 implies the derivative is a linear embedding of patterns at x, and so is 1–quasiconformal. As this holds almost everywhere, work of Gehring implies ϕ is smooth on $F \setminus \Sigma$. By the assumed rigidity of patterns, the derivative at each point is a scalar multiple of one of a finite number of isometric pattern embeddings, where the scalar is bounded between $\frac{1}{K}$ and K. Since Σ cannot disconnect as it is of codimension at least two and the map is C^1 off of Σ, we know that off of Σ we, in fact, have a scalar multiple of a single isometry. Comparing the scaling along different lines in the pattern, we see, possibly after rescaling the metric on Y, that ϕ is isometric on $F \setminus \Sigma$. By continuity, ϕ is therefore an isometry on all of F. This forces the image to be a single flat. A similar use of Gehring's theorem occurs in [4]. As all of our maps are not just quasiconformal but bilipschitz, we are really using an easy special case of this theorem. $\qquad\square$

3.3 Proof of rigidity results

We now prove our main rigidity results from the results of the last section. First, we conclude the analogue of our rigidity results for bilipschitz embeddings of buildings:

Theorem 3.6 *Let X and Y be buildings of equal rank. If there is a bilipschitz embedding of X in Y then there is a linear embedding of their Weyl chamber patterns. If all such pattern embeddings which are K–quasiconformal are conformal, then all K–bilipschitz embeddings of X in Y are, up to rescaling, isometric embeddings of subbuildings.*

Proof The first conclusion is an immediate consequence of Lemma 3.5. For the second, by Proposition 3.4, each flat in X maps by a similarity to a flat in Y. Since the dilation must be equal on two flats intersecting in anything of positive dimension, and all flats can be connected by such a chain of flats, we can rescale the metric on Y so that flats map isometrically to flats. Since every geodesic is contained in a flat, the map is globally isometric and totally geodesic. □

The proof of our main rigidity result requires one really new ingredient, the main technical result of [13].

Proof of Theorem 1.8 Given a quasi-isometric embedding of X_1 into X_2 satisfying the condition of the first bullet point in Theorem 1.8, by passing to the asymptotic cones, we obtain a bilipschitz embedding of buildings. By Theorem 3.6 this is impossible as there are no linear embeddings of their Coxeter systems. If X_1 and X_2 satisfy the conditions of the second bullet point in Theorem 1.8, then the map on asymptotic cones is an isometric embedding by Theorem 3.6 and so, in particular, every flat is mapped to a single flat. By [12, Lemma 7.1.1], the fact that flats go to individual flats in all asymptotic cones implies that every flat maps to within bounded distance (depending only on the quasi-isometry constants) of a single flat. Since intersections of flats encode the Weyl chamber patterns, we know that the Weyl hyperplanes map to within uniformly bounded distance of Weyl hyperplanes. These affine foliations are quite rigid — in particular, by [16, Lemma 7.1.12], the quasi-isometry is at bounded distance from an affine map preserving patterns. Thus, every apartment in the Tits boundary of the domain maps to a well-defined apartment in the Tits boundary of the domain. Further, this maps respects the decomposition into chambers. It is now easy to check that the image is a subbuilding: the isometric embedding of patterns induces an inclusion of Weyl groups $W' < W$, and charts from the domain building structure push forward to charts in the range which are W' compatible and therefore W compatible. By [13, Theorem 3.1], it is the boundary of a subsymmetric space (or building) $Y' \subset Y$, and the image of the quasi-isometric embedding is at bounded distance from Y'. So the

embedding is a quasi-isometry X to Y', and is therefore at bounded distance from an isometry by the main results of either [6] or [12]. (We quote these last results only for simplicity of exposition. At this point, we know much more, eg that the quasi-isometry from X to Y' is bounded distance from an isometry along each flat, so one can use only a small part of the arguments from those papers.) □

4 Pattern embeddings

In this section we do the necessary algebraic analysis to deduce Corollary 1.10 from Theorem 1.8. Since it does not require too much more work, we also analyze essentially all settings in which AN–maps are possible, at least when the range is classical.

The geometric information needed in our rigidity results (and the algebraic information needed for our constructions of exotic embeddings via AN–maps) is closely related to the restricted root systems of semisimple groups over $\mathbb{R}$. Let X be a symmetric space or Euclidean building of rank r. The maximal flats are r–dimensional Euclidean vector spaces (canonically identified with A). Each such flat has a finite collection of affine foliations, along which the flat meets other maximal flats transversely. These are precisely the level sets of the roots.

One result about QI–embeddings says that if X_1 quasi-isometrically embeds in X_2 (of equal rank) then there is a linear isomorphism $A_1 \to A_2$ which preserves these foliations, and that if we know all such linear maps are conformal then the QI–embedding is at bounded distance from isometric. Using the metric on the flats, we can use duality to work in terms of the root vectors rather than the hyperplanes (which are their perpendiculars). In the following, we study pattern preserving maps of flats using this duality.

4.1 Pattern embeddings

Definition 4.1 By a *pattern embedding* of a root system on R_1 in a root system on R_2 we mean a linear embedding of the underlying vector spaces which sends the root vectors of R_1 to scalar multiples of the roots of R_2. We say the pattern embedding is of *equal rank* if the root systems are of equal rank (and so the embedding is induced by an isomorphism of vector spaces), and that the embedding is *conformal* if the linear map is a conformal embedding.

We want to discuss in more detail when a nonconformal, equal rank pattern embeddings exists of one root system into another. We analyze the reducible case in terms of irreducible components, so we begin by studying general pattern embeddings in the irreducible case. Because we are using the dual notion of pattern, our results on pattern embeddings in unequal rank do not say anything about linear embeddings sending affine foliations to affine foliations.

In rank two, which is somewhat special, we are simply discussing patterns of lines in the plane. The irreducible patterns are A_2, BC_2, and G_2 — which are, respectively, three, four and six lines arranged symmetrically — and the reducible pattern $A_1 \times A_1 = D_2$, which is two perpendicular lines. It is straightforward to check that there are pattern embeddings for any pair where the range has at least as many lines as the domain. These rank-two embeddings will be the building blocks for our study of the higher-rank cases. To this end we note that there are three basic nonconformal embeddings: A_2 into BC_2, D_2 into A_2 and D_2 into BC_2 (for this last case there are two distinct embeddings, one conformal and one not). We have excluded the case of target G_2 as it does not occur inside any higher-rank irreducible root system.

In rank $n > 2$ there are three families of irreducible patterns — A_n, BC_n and D_n (this only for $n > 3$ as $D_3 = A_3$ and D_2 is reducible) — plus the exceptional E_6, E_7, E_8 and F_4 (see for example [14] for a detailed description of these exceptional root systems). For reference, the classical root systems are:

(1) A_n for $n \geq 2$ with $V = \{(x_0, \ldots, x_n) \in \mathbb{R}^{n+1} : x_0 + \cdots + x_n = 0\}$ and roots $x_i - x_j$.

(2) B_n for $n \geq 2$ with $V = \mathbb{R}^n$ and roots $\{x_i \pm x_j\} \cup \{\pm x_i\}$.

(3) C_n for $n \geq 2$ with $V = \mathbb{R}^n$ and roots $\{x_i \pm x_j\} \cup \{\pm 2x_i\}$.

(4) BC_n for $n \geq 2$ with $V = \mathbb{R}^n$ and roots $\{x_i \pm x_j\} \cup \{\pm x_i\} \cup \{\pm 2x_i\}$.

(5) D_n for $n \geq 4$ with $V = \mathbb{R}^n$ and roots $\{\pm x_i \pm x_j\}$.

While the root systems B_n, C_n and BC_n are all distinct, they all have the same root patterns. We will refer to this pattern, somewhat sloppily, as BC_n.

Given a root system on V, a two-dimensional subspace $U \subset V$ is *special* if its intersection with the roots is irreducible (so A_2 or BC_2, as G_2 does not occur inside any irreducible root system other than itself). We will call two roots whose span is special *related*. A root system all of whose special subspaces are of type A_2 is called *simply laced* (these are the A, D and E families). As above, simply counting the roots

in various special planes shows that a pattern that embeds in a simply laced system is itself simply laced.

We need some facts about special planes in root systems. We will call a special plane *full* if it is of type A_2 in a simply laced root system or of type BC_2 in doubly laced root system. (All root systems except G_2 which are not simply laced are doubly laced; G_2 is triply laced, but we will not need this, or the usual definitions of these concepts, here.)

Lemma 4.2 *If R is an irreducible root system then:*

- *Any two full planes are equivalent under the action of the Weyl group.*

- *For any two roots r and r' there is a sequence of roots $r = r_0, r_1, \ldots, r_n = r'$ such that the span of r_i and r_{i+1} is a full special plane for all i.*

Proof Let $U \subset V$ be a special plane. Since R is finite, we can choose a linear functional on V which vanishes on U but not on any root outside U. By perturbing this functional we can find a functional which is nonzero on all roots and whose values on roots in U are much smaller than its values on any other roots. Using this functional to define the positive roots, we see that the positive roots in U are generated by the positive simple roots in U. Thus, U comes from an edge of the corresponding Dynkin diagram. It is standard that the Weyl group is transitive both on positive systems and on simple edges of the Dynkin diagram spanned by roots of a fixed length [14]. Since any diagram contains at most one doubled edge, the first part of the lemma follows.

For the second part, we first consider the weaker equivalence generated by all special planes. Since two roots which are not orthogonal necessarily span a special plane, the equivalence classes give a decomposition of V into orthogonal factors, with every root contained in a factor. Thus, irreducibility implies there is only one factor.

This completes the second claim for simply laced root systems. For doubly laced systems we need to see that two roots in a common special plane of type A_2 are connected by a chain of full special planes. The first paragraph above shows that there are at most two Weyl orbits of such planes (one spanned by short roots and one by long roots), and it is trivial to check the claim for one such plane of each type in both BC_n and F_4.   □

Lemma 4.3 *A pattern embedding of irreducible root patterns that maps all full special planes to special planes of the same type is conformal.*

Proof By the earlier observation, this implies the map is conformal on all special planes. Given any two roots in the domain, consider a chain of full special planes connecting them, as in Lemma 4.2. As any successive pair intersect along a line, they all are scaled by the same conformal constant. By irreducibility, the map thus scales all root lines by the same factor. The metric on V is determined by its restrictions to the root lines (this is a linear algebra exercise for A_n and D_n, and all the other root systems are supersets of these) so the map scales the metric by a single constant factor. □

Corollary 4.4 *Every pattern automorphism of an irreducible pattern is conformal.*

This can also be shown directly: After composing with an element of the Weyl group, the automorphism preserves the positive cone, and therefore the simple roots. Since it must also preserve the types of the special plane, this means it corresponds to an automorphism of the Dynkin diagram (ignoring the direction of multiple edges). These are isometric automorphisms of the root systems if the directions are preserved (automatic except for BC_2, G_2 and F_4, which have a conformal pattern automorphism switching the long and short roots).

Lemma 4.5 *Let P and P' be irreducible root systems of equal rank (excluding G_2). If P and P' are both simply laced or neither is, then every embedding of P in P' is conformal.*

Proof If both are simply laced, this is an immediate corollary of the previous lemma. The argument when neither are is the same once one notes that for non-simply laced root systems, irreducibility implies that any two roots are connected by a chain of special planes of type BC_2. □

Conformal embeddings corresponding to inclusions of subroot systems. Up to symmetries, these are all the obvious inclusions of one root system as a subset of another in the coordinates above, which we call *standard* [14]. Each root system includes canonically into the higher-rank root systems in the same family, and $A_{n-1} \subset D_n \subset BC_n$. Here the isomorphism of A_3 and D_3 causes some confusion, as the two inclusions into BC_n for $n > 3$ are distinct.

By the above lemmas, the only possibilities for nonconformal pattern embedding above rank two are one of A, D or E into one of BC or F. We will put off a discussion of the exceptional root systems until later, and first address the classical examples:

Lemma 4.6 *There are, up to the actions of the Weyl groups (and rescaling), two nonconformal pattern embeddings of A_n into BC_m (for $m \geq n$):*

- *The roots $x_i - x_j$ for $i > j > 0$ map to $y_i - y_j$, and the roots $x_i - x_0$ map to y_i.*
- *The roots $x_i - x_j$ for $i > j > 0$ map to $y_i - y_j$, and the roots $x_i - x_0$ map to $y_i + y_1$ (meaning the line $2y_1$ when $i = 1$).*

Proof Since the embedding is nonconformal, the prior lemmas imply that there is some special plane in A_n which maps to one of type BC_2 in BC_n. The Weyl groups are transitive on special planes in the domain and on type BC_2 planes in the range. Thus, we may assume that the plane spanned by $x_1 - x_0$ and $x_2 - x_1$ maps to the plane spanned by y_1 and y_2 in the range. Up to automorphism there are two such maps, depending on whether the missing root line is one of the y_i or $y_1 \pm y_2$. Thus, we may arrange for the map to behave as indicated for the roots $x_i - x_j$ with $2 \geq i > j$.

Assume, for induction, that we can apply an automorphism to put the map in one of the two standard forms for all the roots $x_i - x_j$ with $k \geq i > j$. The root $x_{k+1} - x_k$ is related to all the roots $x_k - x_i$ for $k > i$, so its image must be related to $y_k - y_i$ for all $k > i \geq 1$ and to either y_k or $y_k + y_1$, respectively. It must also be outside the span of those roots because $x_{k+1} - x_k$ is in the domain. The only such roots are $\pm y_k \pm y_s$ for $s > k$. By applying an automorphism of BC_n fixing the first k coordinates, we may assume $s = k + 1$ and that y_{k+1} occurs with positive sign. So $x_{k+1} - x_k$ maps to $\lambda(y_{k+1} \pm y_k)$ with $\lambda > 0$.

Since $(x_{k+1} - x_k) + (x_k - x_i) = x_{k+1} - x_i$ is a root in the domain for all $0 < i < k$, the image must be a scalar multiple of a root. By induction and linearity, the image is $\lambda y_{k+1} + (1 \pm \lambda)y_k - y_i$. Since no roots in the range have three nonzero coefficients and $\lambda > 0$, we must have $x_{k+1} - x_k$ mapping to $y_{k+1} - y_k$, as claimed. $\square$

Lemma 4.7 *Every pattern embedding of D_n (for $n \geq 4$) into BC_m is conformal (and hence standard).*

Proof As before, a nonconformal embedding must send some special plane to one of type BC_2. The Weyl group of D_n is transitive on special planes, so we may assume one of the planes mapping nonconformally occurs in the canonical A_{n-1} inside D_n. By the remark above, this means the pattern embedding can be taken to be one of two maps sending A_{n-1} into $BC_{n-1} \subset BC_n$:

- *The roots $x_i - x_j$ for $i > j > 1$ map to $y_i - y_j$, and the roots $x_i - x_1$ map to y_i (for $i > 1$).*

- The roots $x_i - x_j$ for $i > j > 1$ map to $y_i - y_j$, and the roots $x_i - x_1$ map to $y_i + y_2$ for $i > 1$ (meaning the line $2y_2$ when $i = 2$).

Consider the root $x_1 + x_2$. It is not in the span of A_{n-1} so its image must have nontrivial y_1 component. It must also be related to all $y_i - y_2$ as these are the images of $x_i - x_2$ (for all $i > 2$). The only such roots are $\pm y_1 \pm y_2$. In the first case above, none of these are related to the images of all the $x_i - x_1$ for $i > 2$, so no such pattern embeddings exist. In the second case, the roots $x_3 - x_1$ and $x_3 - x_2$ map to the two roots $y_3 \pm y_2$. Then the roots $x_1 + x_3 = (x_1 + x_2) + (x_3 - x_2)$ and $x_2 + x_3 = (x_1 + x_2) + (x_3 - x_1)$ must both map to roots. Since the y_2 term in the image of $x_1 + x_2$ can only cancel in one of these two sums, the other will have three nonzero coefficients and so will not be a root. $\qquad\square$

For the exceptional root systems the conformal examples are all standard. Since B_4 and C_4 are both isometric subroot systems of F_4, there are nonconformal embeddings (and even AN–maps) from A_4 into F_4. Likewise, there is a unique nonconformal embedding of D_4 into F_4, via a straightforward calculation as above starting with $A_3 \subset D_4$ into B_3 or C_3 inside F_4. For completeness we give it here on a basis (the rest is determined by linearity):

- $x_2 - x_1 \mapsto y_1$.

- $x_3 - x_2 \mapsto y_2$.

- $x_4 - x_3 \mapsto \frac{1}{2}(y_4 + y_3 - y_2 - y_1)$.

- $x_1 + x_2 \mapsto \frac{1}{2}(y_4 + y_3 + y_2 - y_1)$.

This gives a complete picture of pattern embeddings between irreducible root systems. We next turn to the cases where the domain is reducible. For simplicity, we will only discuss the classical root systems here as the other work out similarly. Likewise we will only consider the case of the domain as a product of two factors as those with more can all be built by compositions of such.

Suppose $R = R_1 \oplus R_2$ with each R_i irreducible. To embed R into R' of equal rank, one needs embeddings of both R_i into R' such that the image vector spaces V_i decompose V' as an internal direct sum. We will call the embedding *standard* if it is conformal on both of the components and their images are orthogonal.

The earlier lemmas tell us what the image subspaces can look like:

- A_n includes in D_{n+1} naturally as the root vectors with coordinate sum zero. Composing with signed permutations of the coordinates of D_{n+1} one can get an embedding of A_n with image any hyperplane of the form $\Sigma \pm y_i = 0$.

- The previous embedding can be composed with the inclusions of D_{n+1} into D_m or BC_m for $m > n$. The image subspace is a hyperplane as above in the subspace spanned by $n + 1$ of the coordinates. These are all the conformal embeddings of A_n (for $n > 3$) into these spaces.

- A_n includes into BC_n nonconformally (in two ways), which is of equal rank. This can be composed with the inclusions of BC_n into BC_m for $m > n$. The image subspaces are the span of n coordinates. These are all of the nonconformal embeddings of A_n into BC_m.

- D_n includes into BC_n conformally and of equal rank. This can be composed with any inclusion BC_n into BC_m. The image is the span of n coordinates. These are all the pattern embeddings D_n into BC_m (note that for $n = 3$ this is also an embedding of A_3 into BC_m not covered by the earlier cases).

- For $n \leq m$, A_n includes into A_m, D_n into D_m and BC_n into BC_m. These are conformal with image the span of n coordinates.

The upshot is that the image subspace when the domain is rank n is always either the span of n coordinates or a hyperplane of the specified type in the span of $n + 1$ coordinates, and the latter happens only if the domain pattern is A_n.

Consider an embedding $R_1 \oplus R_2$ into R' of equal rank. If both factors have images of the first type (the span of a subset of coordinates of size equal to the rank of R_i inside D_m or BC_m) then these subsets of coordinates must be complementary to get $V' = V_1 \oplus V_2$. This exhausts the possibilities when no factor is of type A, and in these cases all the embeddings are conformal. It also includes all the cases where A_n only embeds nonconformally, and all such embeddings are simply products of irreducible embeddings composed with one of the conformal embeddings of products of D and BC factors into D or BC.

When the target is of type A, the factors must be so as well. This case is nonconformal:

Lemma 4.8 *There are equal-rank, nonstandard pattern embeddings of $A_m \oplus A_n$ into A_{m+n}.*

Proof Divide the $m + n + 1$ coordinates for A_{m+n} into two sets, one of size $m + 1$ and one of size $n + 1$, which overlap in a single element. The natural inclusions of A_n and A_m as the spans of these are conformal embeddings of A_n and A_m. Their images are disjoint (as the intersection consists of vectors with coordinate sum zero supported in a single coordinate) and together span. $\qquad\square$

Similarly:

Lemma 4.9 *There is a nonstandard pattern embedding of $A_n \oplus A_m$ into D_{m+n}.*

Proof The embeddings of the factors are given by A_k includes in D_{k+1} as the vectors of coordinate sum zero, composed with an isometric embedding of D_{k+1} into D_{m+n}. The image of the underlying vector space is therefore a hyperplane in the span of $k + 1$ coordinates defined by $\sum \pm y_i = 0$ for some collection of signs. If the two sets of coordinates overlap in fewer than two coordinates then the image isn't of full rank. If they overlap in more than two then they are not disjoint. Thus, the only embedding comes from overlapping in exactly two coordinates (say y_1 and y_2) and where the sign choices agree for one of these two and differ for the other. Explicitly, after signed permutations, it must be

$$(x_0, x_1, \ldots, x_n) \times (z_0, z_1, \ldots, z_m) \mapsto (x_0 - z_0, x_1 + z_1, x_2, \ldots, x_n, y_2, \ldots, y_m). \quad \square$$

In the same way we get nonstandard pattern embeddings $A_n \times D_m$ into D_{n+m} and $A_n \times BC_m$ into BC_{n+m} with all factors conformal. Up to compositions, we now have a complete list. To be more precise, we say that a pattern embedding is *maximal* if it only factors as a composition of two embeddings trivially (meaning that one embedding is actually an isomorphism). The above constructions together with the list of possible image subspaces for the irreducible embeddings give a complete list.

Corollary 4.10 *The maximal, equal-rank, nonstandard pattern embeddings among products of the classical root systems (with irreducible range) are:*

- $D_2 = A_1 \times A_1$ *into* A_2.
- A_n *into* BC_n *(in two ways for $n > 2$).*
- $A_n \times A_m$ *into* A_{n+m}.
- $A_n \times A_m$ *into* D_{n+m}.
- $A_n \times D_m$ *into* D_{n+m}.
- $A_n \times BC_m$ *into* BC_{n+m}.

In all of the reducible cases the factors embed conformally but the images are not orthogonal.

Proof   All listed embeddings have been constructed above. It remains to show the list is complete. This is a direct consequence of the classification of irreducible embeddings above and the discussion there of what the image vector spaces can be. First the case when the image is of type A:

The factors are all of type A and embedded conformally. Let the target be A_r, which we view in the standard way as a subset of vectors in $\mathbb{R}^{r+1}$ with coordinate sum zero. Each factor has image the vectors lying in the span of a subset of the coordinates. Two such sets cannot overlap in more than one coordinate or their intersection would contain a root vector. If two overlap in exactly one factor then we have two factors embedding as in the lemma. Thus, our embedding is a composition of two nontrivial embeddings unless these are the only two factors. The last possibility is that all factors land in nonoverlapping sets of coordinates. In this case the span of the images is a subspace of $\mathbb{R}^{r+1}$ of codimension equal to the number of factors. Since we are assuming equal rank, this implies there is only one factor and our map is therefore an isomorphism (and, in particular, conformal).

When the target is D_n or BC_n we argue similarly, but things are simpler as the roots span the full $\mathbb{R}^n$ in the standard coordinates for these root spaces. Again, the image of each irreducible factor is either the span of a subset of the coordinates, or a hyperplane in such. Any two images of the first type must map to disjoint sets of coordinates (as they must be disjoint subspaces). The map thus factors through the map that combines these factors first, hence is not maximal unless there are only the two factors. Since the map is nonconformal while these image spaces are orthogonal, at least one of the factors must be A_k embedding nonconformally into BC_k, but then again the map is nonmaximal.

Thus, there is at most one factor of the first type. Similarly, the argument used above for A_n targets gives that maximality implies there is only one of the second type as well. These are precisely the cases of one A_n factor, the last three on the list.

Finally, suppose there are no factors of the first type and two of the second. This is then an embedding of $A_n \times A_m$ into either D_{n+m} or BC_{n+m}. The former is the third listed embedding. Since the images are assumed of the hyperplane type, the classification of embeddings of A into BC says the image lands in $D \subset BC$, which contradicts maximality.   □

4.2 Relations with AN–maps

Recall that if $A \subset G$ is a maximal $\mathbb{R}$–split torus then diagonalizing the action of A on the Lie algebra of G gives the collection of eigenvalue functions $\lambda_i \colon A \to \mathbb{R}$, which are the restricted roots. If we choose a hyperplane in A to divide the roots into positive and negative, the N in the KAN decomposition is the span of the root spaces $E_\lambda \subset G$ with λ positive. Hence, by equivariance, an AN–map from $A_1 N_1$ into $A_2 N_2$ gives an isomorphism $A_1 \to A_2$ which sends positive roots to positive roots. In particular, it gives a pattern embedding.

A basic Lie theory fact is that if r, r' and $r + r'$ are roots, then, letting E_* be the corresponding root spaces, one has $[E_r, E_{r'}] = E_{r+r'}$. The nontrivial part here is that the bracket is surjective. In the $\mathbb{R}$–split case (or the complex case) this says that any nonzero vectors in E_r and $E_{r'}$ have nonzero bracket. For an AN–map, it follows that the image of the positive roots is *closed*, meaning that if r and r' are roots in the domain then $r + r'$ is a root in the range if and only if it is in the domain. This puts strong limitations on the kinds of nonconformality that can occur.

To make this precise, and extend to the nonsplit case, we focus on the two kinds of basic nonconformality: a D_2 mapping into an A_2 or nonconformally into a BC_2, and an A_2 into a BC_2.

Lemma 4.11 *Let f be an AN–map from $A_1 N_1$ into $A_2 N_2$. If r and r' are positive roots in the domain which generate a D_2 and which map to nonorthogonal roots s and s', then $s + s'$ is not a root.*

Lemma 4.12 *Let f be an AN–map from $A_1 N_1$ into $A_2 N_2$. Suppose P is a special plane for the domain of type A_2 containing positive roots r, r' and $r + r'$ which map to roots s, s' and $s + s'$ in a plane of type BC_2. Neither $s + 2s'$ nor $s' + 2s$ is a root.*

In both cases the conclusion is a restriction on which roots in the image plane are positive, constraining them to lie in a single quadrant. Note that this is the same as saying the image is closed without any assumptions about the sizes of the root spaces (in particular, nothing is assumed $\mathbb{R}$–split).

We need a basic fact from the theory of semisimple Lie algebras (which is also responsible for the fact quoted above in the split case): given a root space E_r and $0 \neq v \in E_r$, there is a vector $u \in E_{-r}$ such that $[v, u] = H_r$, where H_r is the element of the maximal torus such that $[H_r, w] = (r \cdot t)w$ for any root space E_t and any $w \in E_t$.

Proof For the first lemma, let v and v' be nonzero vectors in the root space E_s and $E_{s'}$ which are in the image of f and let u be as above relative to v. Since $r + r'$ is not a root, the images of E_r and $E_{r'}$ commute; in particular, $[v, v'] = 0$. Then, by the Jacobi identity,

$$0 = [[v, v'], u] = [[v, u], v'] + [v, [v', u]] = [H_s, v'] + [v, [v', u]] = (s \cdot s')v' + [v, [v', u]].$$

By assumption, $s \cdot s' \neq 0$, so $[v, [v', u]] \neq 0$, which implies that $[v', u] \neq 0$, so $s' - s$ must be a root. Thus, if $s + s'$ is a root, there is a string of roots $(s' - s)$, $(s' - s) + s$, $(s' - s) + 2s$. This happens only in BC_2 and only if s and s' are orthogonal.

The second lemma is similar.

Let z and z' be elements of the root spaces E_r and $E_{r'}$ with $z'' = [z, z'] \in E_{r+r'}$ nonzero (possible since $[E_r, E_{r'}] = E_{r+r'}$). Let v, v' and $v'' = [v, v']$ be their images in the root spaces E_s, $E_{s'}$ and $E_{s+s'}$, respectively.

Suppose $2s + s'$ is a root in the range (the other case is symmetric). This gives a root string of length three in BC_2: s', $s + s'$, $2s + s'$. As this is the maximal length of a root string, neither $s' - s$ nor $s' + 3s$ is a root in the range (and s and $s + s'$ must be orthogonal). Since $r' + 2r$ is not a root in the domain (the root system A_2 has only length-two root strings), we have $[v'', v] = 0$. Choose a $u \in E_{-s}$ in the range as before. We have, by the Jacobi identity,

$$0 = [0, u] = [[v'', v], u] = [[v'', u], v] + [v'', [v, u]] = [[v'', u], v] - ((s+s') \cdot s)v = [[v'', u], v].$$

The Jacobi identity says $0 = [v'', u] = [[v, v'], u] = [[v, u].v'] + [v, [v', u]]$. The second term on the right is zero as $s' - s$ is not a root, and the first is $(s \cdot s')v'$. Thus, $s \cdot s' = 0$, which is a contradiction as $s \cdot (s + s') = 0$. □

We now turn to discussing the prior pattern embeddings in this context. Since an AN–map must send roots to roots, the two pattern embeddings of A_n into BC_n therefore could come from AN–maps into B_n or C_n, respectively (or either into BC_n).

For the first map, the A_2 planes that get sent to BC_2 planes are the planes P_{ij} spanned by $x_i - x_0$, $x_j - x_0$ and $x_i - x_j$, which map to y_i, y_j and $y_i - y_i$ inside the plane spanned by y_i and y_j. The "missing" roots in the image are $y_i + y_j$ and $-y_i - y_j$. Thus, the lemma above says that y_i and y_j must have opposite signs. For this to be true for all i and j implies that $n = 2$. Thus, this map can arise from an AN–map only for A_2 into BC_2 (in which case this map is the same as the other pattern embedding).

For the second map, the A_2 planes mapping to BC_2 planes are the planes P_i spanned by $x_i - x_0$, $x_i - x_1$ and $x_1 - x_0$, which map to $y_i + y_1$, $y_i - y_1$ and $2y_1$ inside the plane spanned by y_i and y_1. The lemma then says $y_i - y_1$ and $y_i + y_1$ must be the same sign for all i. This happens, for example, if $y_n > \cdots > y_1 > 0$. This map does indeed arise in an AN–map (with this order), as described previously.

In the case of the embedding of D_4 into F_4, it does send roots to roots. However, the image is not closed under addition for any ordering, so there is no AN–map.

For the reducible case, an AN–map for $A_n \times A_m \to A_{n+m}$ was constructed explicitly already, and the $A_n \times D_m \to D_{n+m}$, $A_n \times BC_m \to BC_{n+m}$, etc behave similarly. In all of these cases one can take the Dynkin diagram for the target system and remove one edge, so that the two components are the diagrams for the two factors of the domain. If one orders the roots on the factors so that one is positive and one is negative, additivity is always satisfied.

The one case that does not fit this pattern is $A_n \times A_m \to D_{n+m}$. If we denote the roots in the A_n by $x_i - x_j$ for $n \geq i > j \geq 0$, in the A_m by $z_i - z_j$ for $m \geq i > j \geq 0$, and in the D_{n+m} by $y_i \pm y_j$ for $n + m > i > j \geq 0$, then the map described sends:

- $x_i - x_j$ to $y_i - y_j$.
- $z_i - z_j$ to $y_{i+n-1} - y_{j+n-1}$ for $i > j > 1$.
- $z_i - z_1$ to $y_{i+n-1} - y_1$ for $i > 1$.
- $z_i - z_0$ to $y_{i+n-1} + y_0$ for $i > 1$.
- $z_1 - z_0$ to $y_1 + y_0$.

Since the individual factors map conformally, the only conditions from the lemmas are that the positive roots from one factor must have nonnegative inner product with those from the other. From this we conclude:

- $x_i - x_1$ and $z_j - z_1$ have the same sign for all i and j at least 2.
- $x_i - x_0$ and $z_j - z_0$ have opposite signs for all i and j at least 2.
- $x_1 - x_0$ has opposite sign to both $z_i - z_0$ and $z_i - z_1$ for $i > 1$.
- $z_1 - z_0$ has opposite sign to both $x_i - x_0$ and $x_i - x_1$ for $i > 1$.

This is not possible as $x_2 - x_1$ has the same sign as $z_2 - z_1$, which is opposite to $x_1 - x_0$ and therefore the same as $z_2 - z_0$, hence opposite to $x_2 - x_0$. So $x_2 - x_1$ and $x_2 - x_0$ have opposite signs, which means it is impossible for both to be opposite to $z_1 - z_0$.

Finally, we observe that the closed condition is in general sufficient for the existence of an AN–map when the domain is split, so it is exactly these pattern embeddings that arise from AN–maps. Since the subset is closed, it is immediate that the span of the image root spaces gives a subgroup of N_2 which is graded exactly as N_1 is. Thus, the content of the claim is that this implies the existence of an AN–map. To see this, let AN and AN' correspond to the $\mathbb{R}$–split groups for the given root systems. Let Δ be the set of simple roots in the domain. For each $s \in \Delta$ choose an arbitrary isomorphism $E_s \to E_s'$ between the root spaces in domain and range. The rest is determined inductively. Suppose we already have an isomorphism $E_r \to E_r'$ defined, and $s \in \Delta$ with $r + s$ a root (note that the closed condition says this happens in the domain if and only if it does in the range, so there is no ambiguity). Then $[E_r, E_s] = E_{r+s}$ and $[E_r', E_s'] = E_{r+s}'$, so the map on E_{r+s} is determined. We need to see it is well-defined. If r is a positive root then we need to see that all ways of writing $r = s_1 + s_2 + \cdots + s_n$ with all partial sums roots give the same map $E_r \to E_r'$. To avoid too much Lie theory, we prove this only for the domain A_n, which is sufficient for all the maps constructed here.

Lemma 4.13 *Let s and t be simple roots whose sum is not a root. Suppose r is a positive root such that $r + s$, $r + t$, and $r + s + t$ are all roots. Then the maps $(E_r \otimes E_s) \otimes E_t \to E_{r+s+t}$ and $(E_r \otimes E_t) \otimes E_s \to E_{r+s+t}$ are equal.*

Proof This is essentially just the Jacobi identity. Let u_* be vectors in the relevant root spaces. Then

$$[[u_r, u_s], u_t] = [[u_r, u_t], u_s] + [u_r, [u_s, u_t]].$$

The second term on the right vanishes as E_s and E_t are assumed to commute. $\square$

Thus, we can freely permute adjacent terms provided the simple roots commute. If we express $x_i - x_j$ as above as a sum of simple roots in A_n (namely $x_{s+1} - x_s$ for varying s) then we have the partial sums $x_{i_k} - x_{j_k}$, where at each step either i_k decreases by one or j_k increases by one. The simple roots of those two types always commute, so the lemma allows us to move any one sum to any other.

5 Further questions

We collect here some further questions concerning quasi-isometric embeddings of symmetric spaces left open by our work and which we believe would be a useful guide to further research.

- Can one classify the quasi-isometric embeddings of $\mathrm{Sl}_3(\mathbb{R})$ into $\mathrm{Sp}_4(\mathbb{R})$? It seems unlikely that the AN–maps are the only ones, but we do not know how to construct others. Similarly, can one rule out quasi-isometric embeddings from spaces of A_n to C_n type where no AN–maps exist? Can one rule out embeddings that realize the other linear pattern embedding from Lemma 4.6? Might it be that the existence of a quasi-isometric embedding forces (maybe via a limiting argument) the existence of an AN–map?

- When do quasi-isometric embeddings exist when rank increases? As a start, if X quasi-isometrically embeds in $Y \times \mathbb{R}^d$ for some d, does that imply X quasi-isometrically embeds in Y? More ambitiously, is there any sense in which one can describe all the quasi-isometric embeddings when rank increases? Again, as a start, are all quasi-isometric embeddings of X in $X \times \mathbb{R}$ graphs of Lipschitz functions?

- What can one say about uniformly proper embeddings? These give perhaps a more natural geometric analogue of Margulis's superrigidity. As a warning, recall that the horospheres give a uniformly proper embedding of $\mathbb{R}^n$ in $\mathbb{H}^{n+1}$, so rank is not well behaved.

References

[1] **J Behrstock**, **M F Hagen**, **A Sisto**, *Quasiflats in hierarchically hyperbolic spaces*, preprint (2017) arXiv

[2] **M Bonk**, **O Schramm**, *Embeddings of Gromov hyperbolic spaces*, Geom. Funct. Anal. 10 (2000) 266–306 MR

[3] **N Brady**, **B Farb**, *Filling-invariants at infinity for manifolds of nonpositive curvature*, Trans. Amer. Math. Soc. 350 (1998) 3393–3405 MR

[4] **C Druţu**, *Quasi-isometric classification of non-uniform lattices in semisimple groups of higher rank*, Geom. Funct. Anal. 10 (2000) 327–388 MR

[5] **A Eskin**, *Quasi-isometric rigidity of nonuniform lattices in higher rank symmetric spaces*, J. Amer. Math. Soc. 11 (1998) 321–361 MR

[6] **A Eskin**, **B Farb**, *Quasi-flats and rigidity in higher rank symmetric spaces*, J. Amer. Math. Soc. 10 (1997) 653–692 MR

[7] **B Farb**, *The quasi-isometry classification of lattices in semisimple Lie groups*, Math. Res. Lett. 4 (1997) 705–717 MR

[8] **D Fisher**, **T Nguyen**, *Quasi-isometric embeddings of non-uniform lattices*, preprint (2015) arXiv

[9]　**T Foertsch**, *Bilipschitz embeddings of negative sectional curvature in products of warped product manifolds*, Proc. Amer. Math. Soc. 130 (2002) 2089–2096　MR

[10]　**M Gromov**, *Infinite groups as geometric objects*, from "Proceedings of the International Congress of Mathematicians" (Z Ciesielski, C Olech, editors), volume I, PWN, Warsaw (1984) 385–392　MR

[11]　**J Huang**, *Top-dimensional quasiflats in* CAT(0) *cube complexes*, Geom. Topol. 21 (2017) 2281–2352　MR

[12]　**B Kleiner**, **B Leeb**, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Inst. Hautes Études Sci. Publ. Math. 86 (1997) 115–197　MR

[13]　**B Kleiner**, **B Leeb**, *Rigidity of invariant convex sets in symmetric spaces*, Invent. Math. 163 (2006) 657–676　MR

[14]　**A W Knapp**, *Lie groups beyond an introduction*, 2nd edition, Progress in Mathematics 140, Birkhäuser, Boston (2002)　MR

[15]　**E Leuzinger**, *Corank and asymptotic filling-invariants for symmetric spaces*, Geom. Funct. Anal. 10 (2000) 863–873　MR

[16]　**L Mosher**, **M Sageev**, **K Whyte**, *Quasi-actions on trees, II: Finite depth Bass–Serre trees*, Mem. Amer. Math. Soc. 1008, Amer. Math. Soc., Providence, RI (2011)　MR

[17]　**T Nguyen**, *Quasi-isometric embeddings of symmetric spaces and lattice: the reducible setting*, in preparation

[18]　**P Pansu**, *Métriques de Carnot–Carathéodory et quasiisométries des espaces symétriques de rang un*, Ann. of Math. 129 (1989) 1–60　MR

Department of Mathematics, Indiana University
Bloomington, IN, United States

Department of Mathematics, University of Illinois at Chicago
Chicago, IL, United States

fisherdm@indiana.edu,　kwhyte@math.uic.edu